arXiv:1607.00359v1 [cs.CL] 1 Jul 2016

Moving Toward High Precision Dynamical Modelling in Hidden Markov Models

Sébastien Gagnon, and Jean Rouat PhD,
sebastien.gagnon2@usherbrooke.ca, jean.rouat@usherbrooke.ca *†

July 4, 2016

Abstract

Hidden Markov Model (HMM) is often regarded as the dynamical model of choice in many fields and applications. It is also at the heart of most state-of-the-art speech recognition systems since the 70's. However, from Gaussian mixture models HMMs (GMM-HMM) to deep neural network HMMs (DNN-HMM), the underlying Markovian chain of state-of-the-art models did not changed much. The "left-to-right" topology is mostly always employed because very few other alternatives exist. In this paper, we propose that finely-tuned HMM topologies are essential for precise temporal modelling and that this approach should be investigated in state-of-the-art HMM system. As such, we propose a proof-of-concept framework for learning efficient topologies by pruning down complex generic models. Speech recognition experiments that were conducted indicate that complex time dependencies can be better learned by this approach than with classical "left-to-right" models.

1 Introduction

The correctness of a hidden Markov model's (HMM) topology can strongly influence the model accuracy of a HMM systems, especially for signals with high dynamic variability. This graphical architecture is usually hand-designed to a simple and generic form (usually shared across all classes), whereas constructing precisely tuned class representations can be challenging.

In the 70's a "left-to-right" topology was first proposed for speech modelling, meaning that feature changes through time always flowed in a specific sequential order [1]. It is however simplifying considering that spontaneous speech dynamics are known to be very variable [2]. Up to these days, most state-of-the-art ASR systems such as *deep neural networks*-HMMs (DNN-HMMs, [3]) are still based on that architecture.

As states of HMMs encode static feature space distributions, simple HMM topologies can only model coarse dynamics. A dynamic process constructed from static events is as detailed as the number of such events. Too little precision results in an *underfitted* model with low discriminative power in classification systems. Too much precision, however, can lead to an *overfitted* model without generalizing power [4], making it unable to recognize anything but the training signals. In model selection, as discussed in [4], the key is to balance precision and generalization for maximum performance.

Robustness, the ability of a system to tolerate recording environment changes, is also to be considered and seems to be strongly related to a model's generalizing power [5]. As such, improving model precision would decrease robustness. *Online adaptation* techniques, however, can easily compensate for

*Thanks to S. Brodeur and S. Wood, for fruitful discussions and to J.-F. Duval for his insightful review. Partial Funding by NSERC discovery

†J. Rouat and S. Gagnon are with the NECOTIS research group, dept. of ECE, Univ. de Sherbrooke, Qc., Canada, J1K 2R1

such a drawback. Based on a more comprehensive approach, these methods deal with unseen noise by modifying model statistics at testing time [6, 7, 8]. Performance of such systems in noisy conditions are quite remarkable.

Even if "left-to-right" topologies are good for speech signals, some datasets with higher temporal complexity need more precise architectures. The fact, however, is that topologies are usually hand-designed and kept simple, following Occam's razor principle [9]. Therefore, the main goal of this work is to provide a framework to automatically learn precise HMM topologies from data.

"Left-to-right" HDP-HMM [10], a recently developed technique, is one example of such a framework; it is however unbound by Occam's razor principle or any strong underfitting/overfitting criterion. The expected size of the learned topology being dependant on a *concentration* parameter chosen *a priori*, the designer has indirect control over the degree of temporal precision. Furthermore, the model is allowed to both increase or decrease its size according to a Dirichlet process and is thus unconstrained to follow Occam's principle [9]. On the other hand, this pioneer work allow for the development of HMM speech models without any *a priori* knowledge of the dynamics, something that is not possible with other approaches.

While adept at learning complex topologies from data, "left-to-right" HDP-HMM is oblivious to recognition accuracy and how the architecture influences it. Like most HMM training procedures, this is caused by a mismatch between training and decoding alignement methods, i.e. *forward-backward* vs *Viterbi*. In [10], for example, some learned monophone topologies allow decoded paths to be as short as 2 time frames long whereas in standard monophone systems the shortest path is 3 time frames. According to our preliminary experiments, this tend to generate insertion errors when decoding, enough to significantly lower accuracy as defined by the word error rate (WER) standard. As shown in TIMIT benchmarks listed in [11], considering insertion errors always significantly decreases the recognition performance of a system. Thus, while useful in approximating the topology needed for each class, "left-to-right" HDP-HMM procedure alone is not ideal for our intended goal.

Conventional approaches to dynamics encoding with HMMs usually substitute topology learning with transition probabilities estimation. In speech recognition, this is the most popular paradigm: generic "left-to-right" architectures are adapted to target signal's dynamics by tuning a few persistent parameters. In [12], such an approach is attempted on a complex generic topology with improved additive noise robustness. However, clean speech performance are not reported, which may suggests that precision has not improved. In fact, improved robustness can be linked to a greater generalizing power, as explored in [5].

These results might be explained by an intrinsic problem of HMMs, the imbalance between the dynamic ranges of the transition and emission probabilities. Exposed by Rabiner and Huang in [13], this phenomenon is at the root of the popular thought that transition probabilities are almost useless. It is even a common practice for designers to implement HMM ASR systems with untrained transitions, because the loss in performance is fairly small. Explained in [13] as a lack of pervasive discriminative power of the transition probabilities in path decoding, we conceptualize its effects as rendering equiprobable all transitions that leave the same state. Thus, tuning transition probabilities cannot be a good substitution to complex topology learning.

In this work, we first analyse the effects of the imbalance phenomenon. We show that all paths leaving the same state are effectively equiprobable in the standard TIMIT monophone recognition experiment. Thus, topology learning is shown essential for precise dynamics modelling, for which we then propose a simple and accessible framework. Assuming that HMM spoken word models in conventional ASR systems are closer to underfitting than overfitting (a reasoning we based on [2]), we propose to use model *flattening* [12] in conjunction with *transitions pruning* to extract precise class topologies. *Flattening* is the process of transforming a simple "left-to-right" *Gaussian mixture model*-HMM (GMM-HMM) into an equivalent complex HMM with single Gaussian emission models. Using *transitions pruning* to reduce the flatten model complexity then reveals a more precise dy-

namic model while still following Occam's razor principle. We finally demonstrate that with the same number of emission model parameters, our technique clearly outperforms the classic "left-to-right" topology on clean word recognition tasks.

2 Transition and Emission Probabilities Imbalance

As discussed in [13], transition probabilities may not play a significant role in path decoding (using Viterbi); recognition could be entirely independent of them. Then, all transitions leaving the same state could be considered effectively equiprobable during path decoding. To the knowledge of the authors, this phenomenon has not been quantitatively documented for speech recognition. In the *token passing* implementation of the Viterbi algorithm, the state $s(t+1)$ occupied at time $t+1$ is given by [14]:

$$s(t+1) = \underset{k=1,2,...,N}{\operatorname{argmax}} [a_{s(t)\to k} b_k(O_{t+1})] \quad (1)$$

Where N is the total number of emitting states in the model. Let there be a distinction between zero and non-zero transition probabilities:

$$\forall (i,k) \in [1,...,N]; a_{i\to k} \neq 0 \to k \in \varphi_i \quad (2)$$

Were φ_i regroups all non-zero transitions leaving state i. Thus, (1) can be reformulated in the following fashion:

$$s(t+1) = \underset{k\in\varphi_{s(t)}}{\operatorname{argmax}} [a_{s(t)\to k} b_k(O_{t+1})] \quad (3)$$

Equation (3) only takes into account states that are linked by a non-zero transition from $s(t)$. Formula (3) can be formulated as follows:

$$s(t+1) = i \;\; if \;\; \frac{b_i(O_{t+1})}{b_k(O_{t+1})} > \frac{a_{s(t)\to k}}{a_{s(t)\to i}}, \forall (i,k) \in \varphi_{s(t)}, k \neq i \quad (4)$$

$$\text{Lets define: } \beta = \frac{b_i(O_{t+1})}{b_k(O_{t+1})} \;\; \alpha = \frac{a_{s(t)\to k}}{a_{s(t)\to i}}$$

We defined α and β as ratios of probabilities to isolate the respective transition and emission discriminability forces. The variance of these variables, respectively the transition and emission discriminability coefficients, give a good estimate of their dynamic ranges. To evaluate them, we conducted an experiment on the TIMIT training set with conventional 5-states (3 emitting states) "left-to-right" monophonic models with Gaussian mixture models (GMMs) of 16 components on each state. The procedure is done in 2 steps for each training utterance: first, using the appropriate models (listed in the signal's label) an ideal path is computed with the *forward-backward* algorithm [15]. Then, for each transition taken in the decoded path the α and β values are computed. The variances are then estimated across all the training set. The path alignment method used here, forward-backward, is not equivalent to (4). In fact, Token passing does not take the backward probability into account and is therefore less optimal. This was done purposefully to favor high emission probabilities in an effort to minimize discriminative power. One must understand that strong model mismatch comes from emission probability values being several orders of magnitude different from one state to another, which far more happens in low probabilities (in mismatched dynamics). In other words, the emission discriminability coefficient calculated is minimized to a level unattainable in practical applications.

$$\sigma(ln(\alpha)) = 0.80 \;\; \sigma(ln(\beta)) = 193.24$$

Where $\sigma(i)$ is the variance of i. In the linear domain, the standard deviation of β is roughly 440,000 times larger than α. Thus, the transition probabilities are in some sense binary variables, i.e. they are, or not, members of $\varphi_{s(t)}$ in (4).

This is because emission probabilities have a near-infinite dynamic range, while transition probability do not, for any given $s(t)$. In fact, considering how this problem is exposed in [13], we infer that only in topologies with states of near-infinite branching ratios may this imbalance vanishes. It is therefore safe to assume that all discrete topologies considered in this work are equally affected by this imbalance.

3 Pruning

Encoding acoustic dynamic properties in a generic HMM model is to change its topology, i.e. by activating or deactivating transitions. A deactivated transition has a probability value of 0 and is therefore not involved in (4).

Learning the topology can be done in 3 ways: either "growing" from a simple prototype model (ex. [16]), "pruning" from a complex generic model (ex. [17]) or a mixture of both (ex. [10]). With "growing" techniques, i.e. increasing the model's complexity, an almost groundless guess must be made to determine how the expansion is done. This is very much subject to human error.

On the other hand, "pruning" processes are much more reliable as one removes only the paths that are not often visited. Mak and Chan [17], for example, have successfully used pruning on a "left-to-right" topology with long range transitions. When compared with an unpruned system, they obtained a significant improvement on the accuracy in a clean word recognition task. Our work follows that line of thinking and implements pruning instead of other alternatives.

4 Proposed System

4.1 Integration of the Pruning Module and Threshold Optimization

A modification of the standard HMM training procedure is proposed for increasing the temporal modelling precision. Fig. 1 illustrates the full proposed system. The pruning (step #5) is done by comparing each individual transition probability with a threshold value, ϵ:

$$\text{if } a_{i \to j} > \epsilon \text{ then keep, else } a_{i \to j} = 0 \qquad (5)$$

Since the value of ϵ is unknown, optimization steps are required during training to find ϵ (using the loop step #6 in Fig. 1). The very simple ϵ optimization process is designed as such to exploit the steady relation between pruning threshold and performance, thus avoiding unpredictable local minimums.

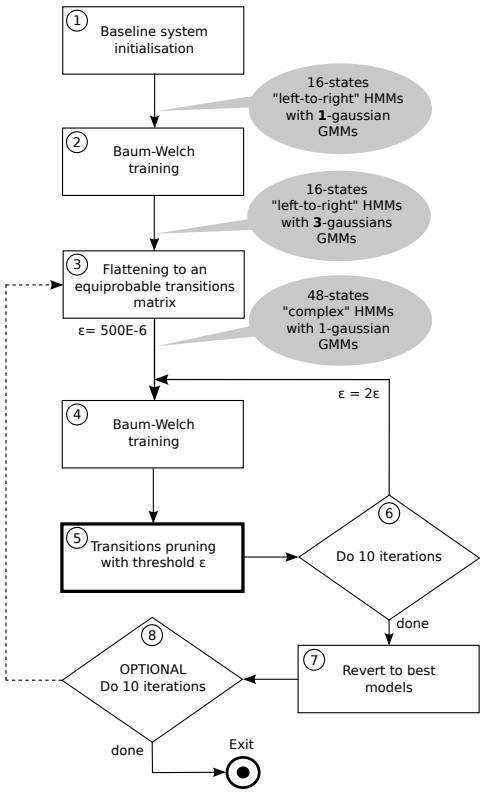

Figure 1: 1) Flat initialisation; 2) Conventional Baum-Welch training with GMM mixture splitting; 3) Flattening process on the baseline GMM-HMM, transition matrix is equiprobable on the allowed transitions; 4) Baum-Welch training without GMM mixture splitting; 5) Pruning as per (5); 7) The model with highest decoding accuracy on the training set is kept, all the others are discarded; 8) (Optional) Emission models feedback: emission models of the model kept in 7 are given to their respective states in the flatten model in 3. All pruned transitions are thus reactivated.

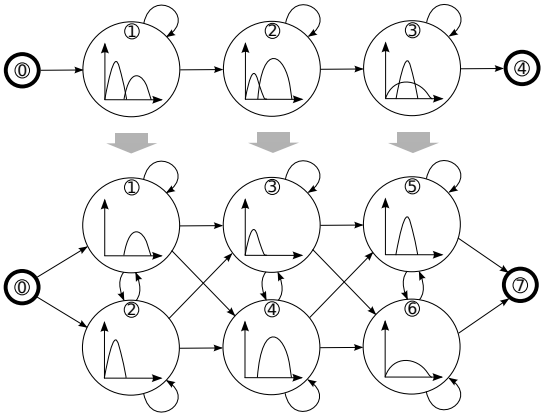

Figure 2: Flattening process; the top "left-to-right" model has a 2 Gaussian mixture model for each state; the bottom "complex" model is a flatten version of the top model with only 1 Gaussian per state.

4.2 Initialization by Model Flattening

To maximize the beneficial impact of transition pruning we work on a complex prototype model (high number of states and transitions). However, this can be difficult since the transition parameter space is larger than with simple "left-to-right" topologies. Furthermore, if it is improper, alternate paths tend to die off during training (i.e. very low occupancy probability) to only favour a single path through the topology. This effectively returns the model to a long "left-to-right" chain with mediocre performance. Thus, the initialization of a complex HMM model is an important aspect of this work, for which the *flattening* technique presented in [12] is used. Since a multi-gaussian mixture model can be viewed as an HMM of single-gaussian states, the flattening consists in replacing the states of a "left-to-right" GMM-HMM by their respective HMM's form. This effectively flattens the representation to a lattice of mono-gaussian densities, as illustrated in Fig. 2. The final form of a trained "left-to-right" model is already fine-tuned to the acoustic properties of the modelled class, it is therefore an ideal configuration.

4.3 Feedback of Emission Models

To further increase the recognition performance of the proposed system, the emission models of the pruned HMMs are fed back to the initialization step (link between steps #8 and #3 in Fig.1):

$$\forall i \in [1, ..., N]; b_i^{pre}(\cdot) \leftarrow b_i^{post}(\cdot) \tag{6}$$

Where $b_i^{pre}(\cdot)$ and $b_i^{post}(\cdot)$ are respectively the emission distributions of the state i at initialization and after training. As complex models require more care to train adequately because of their vast transition parameter space, this step is added to allow slower convergence. We observed experimentally that with about 10 iterations of this "feedback" process (step #8 in Fig. 1) the recognition performance on the training set seems to saturate. While this step is not required to beat baseline accuracies on clean word recognition, on average it increases performance even further (Table 1). However, it also sharply increases the computational effort required for training.

5 Experimental Framework

To evaluate the temporal modelling precision, we chose to view it in terms of recognition accuracy. Our premise is that classical "left-to-right" modelling of speech dynamics is closer to *underfitting* than *overfitting* and therefore improved precision should yield better results. According to [2], this is most likely true for spoken word models. Furthermore, clean speech accuracies (the tested signal was recorded in the same conditions as the training signals) are the main focus since, as explored in [5], noisy recognition seems to deal more with generalizing power than precision.

For all conducted tests, we want any increase in recognition accuracy to be entirely attributed to the higher temporal modelling precision. This is done by ensuring that the total amount of Gaussian distributions for each HMM was the same, i.e. the summed amount of GMM mixture components across each model is identical for both the reference and the proposed systems.

First, large dictionary speech recognition applicability is evaluated with a word recognition task with

monophonic models on TIMIT. The baseline "left-to-right" 5-states (3 emitting states) models have 10 Gaussian mixture components per state. The proposed system uses flatten versions of the "left-to-right" models and are thus 30 states long with single-Gaussian GMMs.

Next we tested on the Aurora-2 digits word classification task. Since monophonic models should not posses much temporal structure, by linguistic definition, a word recognition task is thought to be more representative of the difference in temporal modelling precision (word models encode much richer dynamics than monophones). The baseline "left-to-right" HMM models used for the Aurora-2 tasks are 18-states long (16 emitting states) with 3-Gaussian GMMs. The proposed models are 50-states long (48 emitting states) with 1-Gaussian GMM per state. Training is done on the Aurora-2 clean speech "TRAIN" corpus. Noisy recognition tasks are also performed to evaluate the robustness of the technique.

Since the clean speech recognition accuracies of the reference "left-to-right" GMM-HMMs on this last dataset are very high (>99%), improvements may not be significant. As a result, we generated 4 new versions of the Aurora-2 digits dataset, each of them convoluted with a different real world reverberation impulse response (IR) taken from the Openair IR database [18]. The IRs are chosen on the basis of their uniqueness. We selected "Maes Howe", "Falkland Palace Royal Tennis Court", "Purnode's Tunnel" and "Tyndall Bruce Monument". For these experiments, both systems are trained on a convoluted version of the "TRAIN" corpus. Clean and noisy testing datasets are also convoluted in the same fashion, which means the latter is corrupted by IRs with its additive noise.

Word error rates (WER) are computed in the standard fashion, taking into account insertion and deletion errors. For every one of the 3 Aurora-2 sets (A, B, and C), presented WER values are averaged over all additive noise types. Furthermore, noisy tests are averaged over SNRs 20, 15, 10, 5 and 0 dB.

Table 1: Performances measured on the Aurora-2 and TIMIT datasets

		WER (%)			Relative WER reduction vs baseline (%)	
AURORA-2						
Convolutive noise	**Additive noise SNR**	**Baseline**	**Proposed, no feedback**	**Proposed, 10 iterations feedback**	**Proposed, no feedback**	**Proposed, 10 iterations feedback**
Clean	Clean	0.93	0.81	0.70	12.59	24.97
	20dB-0dB	40.33	47.06	45.75	-16.71	-13.46
Maes Howe	Clean	6.54	5.67	5.70	13.43	12.93
	20dB-0dB	55.42	55.32	55.96	0.17	-0.97
Tennis court	Clean	56.88	65.30	60.70	-14.81	-6.71
	20dB-0dB	78.91	91.64	90.25	-16.13	-14.38
Tunnel	Clean	27.18	25.50	20.16	6.19	25.82
	20dB-0dB	69.72	77.54	73.55	-11.22	-5.49
Tyndall Bruce	Clean	18.52	16.32	17.08	11.85	7.77
	20dB-0dB	58.91	63.80	67.35	-8.29	-14.31
Average	Clean	22.01	22.72	20.87	5.85	12.96
	20dB-0dB	60.66	67.07	66.57	-10.44	-9.72
TIMIT						
Clean	Clean	31.98	-	35.88	-	-12.20

6 Results and Discussion

Table 1 shows significant WER reductions with the proposed approach in clean word recognition. When trained and tested in reverberated environments without additive noise, performance are also significantly improved on average. Since reverberation adds time-dependencies to signals and thus increases their temporal complexity, this demonstrates that temporal precision has indeed been increased. However, the proposed system is also less robust to additive noise. This may be explained by a loss of generalizing power caused by the improved precision [5]. To our knowledge, the poor performance with the TIMIT database (shown in Table 2) are best explained by the fact that monophonic HMMs have linguistically little to no temporal structure, as the name implies. Hence, a model specially designed to encode complex temporal behaviours is unsuited to this recognition task.

We thus see that monophone-based speech recognition does not seem to profit from models with higher temporal precision while word-based classification does. As discussed in [2], this indicates that a high temporal variability exists at the syllable level. Considering this, we suggest that the proposed approach should perform better in large dictionnary systems with classes representing linguistic units of increased length, such as triphones [11].

The proposed proof-of-concept framework could be

improved in a number of ways for hypothetical increased performances. First, a less trivial optimization algorithm (see Fig.1) could be used in the pruning iterative mechanism. This could yield very precise pruning strengths leading to high recognition accuracy.

Better initialization models and methods could also greatly benefit our solution. As it was discussed earlier, complex and finely-tuned models are very sensitive to their initialization conditions and as such, there is much work to be done in optimizing them. Finally, the proposed framework can also be coupled with complementary state-of-the-art techniques that implement better emission models such as DNN-HMMs [3] and online model adaptation [7].

References

[1] F. Jelinek, "Continuous speech recognition by statistical methods," *Proceedings of the IEEE*, vol. 64, no. 4, pp. 532 – 56, 1976/04.

[2] S. Greenberg, "Speaking in shorthand-a syllable-centric perspective for understanding pronunciation variation," *Speech Communication*, vol. 29, no. 2-4, pp. 159 – 76, 1999/11.

[3] A. Mohamed, G.E. Dahl, and G. Hinton, "Acoustic modeling using deep belief networks," *IEEE Transactions on Audio, Speech and Language Processing*, vol. 20, no. 1, pp. 14 – 22, 2012/01.

[4] Gavin C. Cawley and Nicola L. C. Talbot, "On over-fitting in model selection and subsequent selection bias in performance evaluation," *Journal of Machine Learning Research*, vol. 11, pp. 2079 – 2107, 2010.

[5] Xiong Xiao, Jinyu Li, Eng Siong Chng, Haizhou Li, and Chin-Hui Lee, "A study on the generalization capability of acoustic models for robust speech recognition," *IEEE Transactions on Audio, Speech and Language Processing*, vol. 18, no. 6, pp. 1158 – 69, 2010/08.

[6] Ozlem Kalinli, Michael L. Seltzer, Jasha Droppo, and Alex Acero, "Noise adaptive training for robust automatic speech recognition," *IEEE Transactions on Audio, Speech and Language Processing*, vol. 18, no. 8, pp. 1889 – 1901, 2010.

[7] Jinyu Li, Li Deng, Dong Yu, Yifan Gong, and A. Acero, "High-performance HMM adaptation with joint compensation of additive and convolutive distortions via vector taylor series," in *2007 IEEE Workshop on Automatic Speech Recognition and Understanding*, Piscataway, NJ, USA, 2007//, pp. 65 – 70.

[8] A. Narayanan and DeLiang Wang, "Improving robustness of deep neural network acoustic models via speech separation and joint adaptive training," *IEEE/ACM Transactions on Audio, Speech and Language Processing*, vol. 23, no. 1, pp. 92 – 101, 2015/01.

[9] David J.C. MacKay, *Model Comparison and Occam's Razor*, pp. 345 – 357, Cambridge, UK, 2003.

[10] A.H.H.N. Torbati, J. Picone, and M. Sobel, "A left-to-right HDP-HMM with HDPM emissions," in *2014 48th Annual Conference on Information Sciences and Systems (CISS)*, Piscataway, NJ, USA, 2014, p. 6 pp.

[11] C. Lopes and F. Perdigao, *Phone Recognition on the TIMIT Database*, pp. 285 – 302, Rijeka, Croatia, 2011.

[12] Yong Zhao and Biing-Hwang Juang, "Stranded gaussian mixture hidden markov models for robust speech recognition," in *Proceedings of the 2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2012)*, Piscataway, NJ, USA, 2012, pp. 4301 – 4.

[13] L. R. Rabiner and B. H. Juang, "Hidden markov models for speech recognition - strengths and limitations," *Speech Recognition and Understanding. Recent Advances, Trends and Applica-*

tions. Proceedings of the NATO Advanced Study Institute, pp. 3 – 29, 1992.

[14] S.J. Young, N.H. Russell, and J.H.S. Thornton, "Token passing: a simple conceptual model for connected speech recognition systems," 1989.

[15] L.R. Rabiner, "A tutorial on hidden markov models and selected applications in speech recognition," *Proceedings of the IEEE*, vol. 77, no. 2, pp. 257 – 86, 1989/02/.

[16] T. Jitsuhiro and S. Nakamura, "Automatic generation of non-uniform hmm structures based on variational bayesian approach," in *2004 IEEE International Conference on Acoustics, Speech, and Signal Processing*, Piscataway, NJ, USA, 2004, vol. vol.1, pp. 805 – 8.

[17] B. Mak and Kin-Wah Chan, "Pruning hidden markov models with optimal brain surgeon," *IEEE Transactions on Speech and Audio Processing*, vol. 13, no. 5, pp. 993 – 1003, Sept. 2005.

[18] D. Murphy and S. Shelley, "Open acoustic impulse response (open air) library," .